

APPLICATION FOR UNITED STATES LETTERS PATENT

FOR

COMPUTING SYSTEM CAPABLE OF REDUCING POWER CONSUMPTION BY
DISTRIBUTING EXECUTION OF INSTRUCTION ACROSS MULTIPLE PROCESSORS
AND METHOD THEREFORE

Inventor(s): Matthew A. Hayduk

Prepared by: Kenneth M. Seddon,
Senior Patent Attorney

intel®

Intel Corporation
5000 W. Chandler Blvd., CH6-404
Chandler, AZ 85226-3699
Phone: (480) 554-9732
Facsimile: (480) 554-7738

"Express Mail" label number **ET 616076811 US**

COMPUTING SYSTEM CAPABLE OF REDUCING POWER CONSUMPTION BY
DISTRIBUTING EXECUTION OF INSTRUCTION ACROSS MULTIPLE PROCESSORS
AND METHOD THEREFORE

5 BACKGROUND

To increase the overall performance of some computing systems, it may be desirable to design the computing system so that it includes two or more processors. The use of multiple processors typically increases the total available instruction cycles that may be used to execute instructions on the computing system. The number of instructions cycles available with a processor is typically measured in millions of instructions per second (MIPS), and it generally follows that the more MIPS associated with a computing system, the faster the computing system may be able to execute multiple programs (e.g. application programs, operation system programs, etc.)

However, the addition of multiple processors to a computing system may significantly increase the total amount of power consumed by the computing system. Such an increase is typically not desirable in mobile computing systems as this may directly reduce the total available computing time due to the increase in the drain on the batteries.

Thus, there is a continuing need for better ways to reduce the power consumption of multiple-processor computing systems while providing increased execution capacity.

BRIEF DESCRIPTION OF THE DRAWINGS

The subject matter regarded as the invention is particularly pointed out and distinctly claimed in the concluding portion of the specification. The invention, however, both as to organization and method of operation, together with objects, features, and advantages thereof, may best be understood by reference to the following detailed description when read with the accompanying drawing in which:

the sole figure is a block diagram representation of a portion of a computing system in accordance with an embodiment of the present invention.

It will be appreciated that for simplicity and clarity of illustration, elements illustrated in the figure have not necessarily been drawn to scale. For example, the dimensions of some of the elements are exaggerated relative to other elements for clarity.

DETAILED DESCRIPTION

In the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the invention. However, it will be understood by those skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known methods, procedures, components and circuits have not been described in detail so as not to obscure the present invention.

Some portions of the detailed description that follows are presented in terms

of algorithms and symbolic representations of operations on data bits or binary digital signals within a computer memory. These algorithmic descriptions and representations may be the techniques used by those skilled in the data processing arts to convey the substance of their work to others skilled in the art.

5 An algorithm is here, and generally, considered to be a self-consistent sequence of acts or operations leading to a desired result. These include physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven
10 convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers or the like. It should be understood, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities.

15 Unless specifically stated otherwise, as apparent from the following discussions, it is appreciated that throughout the specification discussions utilizing terms such as "processing," "computing," "calculating," "determining," or the like, refer to the action and/or processes of a computer or computing system, or similar electronic computing device, that manipulate and/or transform data represented as
20 physical, such as electronic, quantities within the computing system's registers and/or memories into other data similarly represented as physical quantities within the computing system's memories, registers or other such information storage,

transmission or display devices.

Embodiments of the present invention may include apparatuses for performing the operations herein. This apparatus may be specially constructed for the desired purposes, or it may comprise a general purpose computing device

5 selectively activated or reconfigured by a program stored in the device. Such a program may be stored on a storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), electrically programmable read-only memories (EPROMs), electrically erasable and

10 programmable read only memories (EEPROMs), magnetic or optical cards, or any other type of media suitable for storing electronic instructions, and capable of being coupled to a system bus for a computing device.

The processes and displays presented herein are not inherently related to any particular computing device or other apparatus. Various general purpose systems

15 may be used with programs in accordance with the teachings herein, or it may prove convenient to construct a more specialized apparatus to perform the desired method. The desired structure for a variety of these systems will appear from the description below. In addition, embodiments of the present invention are not described with reference to any particular programming language. It will be

20 appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

In the following description and claims, the terms "coupled" and "connected,"

along with their derivatives, may be used. It should be understood that these terms are not intended as synonyms for each other. Rather, in particular embodiments, “connected” may be used to indicate that two or more elements are in direct physical or electrical contact with each other. “Coupled” may mean that two or more elements are in direct physical or electrical contact. However, “coupled” may also mean that two or more elements are not in direct contact with each other, but yet still co-operate or interact with each other.

Turning to FIG. 1, an embodiment 100 in accordance with the present invention is described. Embodiment 100 may comprise a portable communication device 50 such as a mobile communication device (e.g., cell phone), a two-way radio communication system, a one-way pager, a two-way pager, a personal communication system (PCS), a portable computer, or the like. Although it should be understood that the scope and application of the present invention is in no way limited to these examples. Other embodiments of the present invention may include other computing systems that may or may not be portable or even involve communication systems such as, for example, desktop or portable computers, servers, network switching equipment, etc.

In this particular embodiment portable communication device 50 may include two processors 10 and 20 that may execute instructions such as instructions stored in a memory 40. Processors 10 and 20 may be one of a variety of integrated circuits such as, for example, a microprocessor, a central processing unit (CPU), a digital signal processor, a microcontroller, a reduced instruction set computer (RISC), a complex

instruction set computer (CISC), or the like, although the scope of the present invention is not limited by the particular design or functionality performed by processors 10 and 20. In addition, in some embodiments, processors 10 and 20 may be the same processor, while in other embodiments they may be different. For example, in this particular embodiment processor 10 may comprise a CISC processor to execute general user applications while processor 20 may comprise a base band processor that may be used to initiate and receive wireless communications.

Memory 40 may comprise any variety of volatile or non-volatile memory such as any of the types of storage media recited earlier, although this list is certainly not meant to be exhaustive and the scope of the present invention is not limited in this respect. In this particular application, memory 40 may comprise a variety of instructions to be executed by processors 10 or 20. Buses 41 and 42 may be used to allow processors 10 and 20, respectively, to access or store information and/or instructions stored in memory 40.

For example, memory 40 may be used to store sets of instructions such as instructions associated with an application program, an operating system program, a communication protocol program, etc. For example, the instructions stored in memory 40 may be used to perform wireless communications, provide security functionality for portable communication device 50, user functionality such as calendaring, email, internet browsing, telephony services such as voice calls, SMS, call management, fax, or multimedia services such as MP3, speech recognition, voice memo pad, audio codecs, etc.

Processors 10 and 20 may comprise monitors 11 and 21, respectively that may be used to determine, track, or predict the current operational characteristics of processors 10 and 20. For example, monitors 11 and 21 may comprise any combination of hardware and circuitry to determine the performance capability and consumption of processors 10 and 20. Simply stated, monitors 11 and 21 may be used to determine which of processors 10 and 20 is best suited to execute a set of instructions so as to take into account the amount of power consumed by portable communication device 50 and/or the desired performance characteristics of a user.

One technique to determine the performance capability of processors 10 and 20, although the scope of the present invention is not limited in this respect, is that monitors 11 and 21 may track or calculate the number of instructions that may be executed (e.g. the number of millions of instructions per second (MIPS) that may be executed.). Such operational characteristics may be determined by tracking or monitoring one or more values within processors 10 and 20. For example, monitors 11 and 21 may monitor such values as the operational voltage potential (e.g. the voltage potential of a portion of a core within the processor), operational frequency (e.g. one of the internal synchronizing signals), number of wait states generated, the rate at which memory external to the processors is accessed, cache hit/miss rates, amount of current passing through all or a portion of the processors 10 or 20. Although it should be understood that this list is not exhaustive and alternative embodiments may use any combination of these characteristics or others.

Monitors 11 and 21 may comprise any combination of hardware and software.

Hardware may include, although not necessarily limited to, latches, counters, and registers to record the number of instructions cycles or the number of instructions being executed by the corresponding processor. Although the scope of the present invention is not limited so as to involve any particular hardware or any hardware at all as alternative embodiments may implement monitors 11 and 21 using software.

In alternative embodiments, monitors 11 and 21 may be implemented all or in part through software such as sets of instructions executed on processors 10 or 20. As shown in the figure, processors 10 and 20 may comprise memory 13 and 23, which may include any of the types of memory recited earlier. Memory 13 and 23 may be used to store data to track the operational performance of the corresponding processor 10 or 20. For example, memory 13 and 23 may comprise registers or be capable of storing table(s) of data that are used to track the historical average that processors 10 and 20 have taken to execute particular sets of instructions. These tables may also be dynamically updated or averaged over time to so that the values in the table(s) represent more recent information on the overall resource consumption that processors 10 and 20 take to execute particular sets of instructions.

The table(s) may also keep track of the performance characteristics of processors 10 and 20 under various operational parameters. For example, although the scope of the present invention is not limited in this respect, the tables make keep track of the performance characteristics depending on such factors as the

operational voltage, clock frequency, or depending on how many other instructions are currently being executed in processors 10 and 20 or within portable communication device 50.

Although the scope of the present invention is not limited in this respect, portable communication device 50 may also be adapted to store user parameters that may, among other things, allow a user to define or set the desired performance level(s) for executing particular instructions. Accordingly, when portable communication device 50 is to execute a new set of instructions (e.g. an user application in memory 40), it is first determined which of processors 10 and 20 is best suited to execute the instructions. An example will now be provided to illustrate how this determination may be made, although it should be understood that the scope of the present invention is not limited to this example.

In this embodiment, processors 10 and 20 are already executing other programs when a request is made to execute a set of instructions in memory 40 using processor 10. Monitor 11 may first determine the current performance characteristics of processor 10 using one or more of the techniques described earlier. Monitor 11 may then determine if processor 10 may execute the new set of instructions within the desired timeframe acceptable to the user. If processor 10 has sufficient bandwidth (i.e., may execute the set of instructions within acceptable parameters), then the set of instructions may be executed by processor 10.

On the other hand, if processor 10 cannot execute the set of instructions

under the current operating parameters within sufficient time, then monitor 11 may poll the other processors (e.g. processor 20), to determine if the set of instructions may be executed by another processor. Although the scope of the present invention is not limited in this respect, monitor 11 may use a bus 70 to

5 communicate with or consult with monitor 21 to determine if the set of instructions may be executed on processor 20. For example, bus 70 may be a serial or parallel bus that allows monitors 11 and 21 to share performance information about their respective processors 10 and 20. It should be understood that the scope of the present invention is not limited to any particular embodiment for how monitors 11
10 and 21 consult with each other. In alternative embodiments, this may be done at a layer three network level or at the link level. Alternatively, the consultation may take place through the use of interrupts and/or sharing data through memory 40.

Simply stated, monitor 11 consults with monitor 21 to determine if the set of instructions may be executed by processor 20 within the desired performance
15 parameters. If processor 20 has sufficient bandwidth (i.e. sufficient capacity), responsibility for the execution of the set of instructions may be transferred from processor 10 to processor 20.

In yet another embodiment of the present invention, monitors 11 and 21 may consult to determine which processor is better suited to execute the set of
20 instructions to either reduce the overall power consumption of portable communication device 50 or to execute the set of instructions in the least of amount of time. For example, even if processor 10 has sufficient capability to

execute the set of instructions, it may be desirable to have the instructions executed by processor 20 instead while reducing the power consumption associated with processor 10 (e.g. reduce the operational voltage potential, reducing the clock frequency, etc.). Consequently, the set of instructions may get
5 executed within an acceptable time frame while reducing the overall power consumption of portable communication device 50.

If monitor 21 indicates that processor 20 may not have the bandwidth to be able to execute the set of instructions within an acceptable time frame, monitor 11 may call for an increase in the operational parameters (e.g. operating voltage, clock
10 frequency, etc.) of processor 10 that may increase the overall performance capability of processor 10. The increase in operational parameters may increase the MIPS available on processor 10 so that the set of instructions may be executed within an acceptable performance level. It should be understood that processors
10 and 20 may include circuitry that may be used to raise and lower operational
15 parameters such as voltage potential, clock frequency, etc.

In yet other alternative embodiments, algorithms may be used to adjust the operational parameters of processors 10 and/or 20 to balance the computing needs of portable communications device 50 with the overall power consumption of the processors within portable communication device 50. In other embodiments
20 monitors 11 and 21 may be used to balance the programs and/or instructions that are executed by processors 10 and 20 so as to reduce the overall power consumption of portable communication device 50. Additionally, monitors 11 and

21 may be used to control the operational parameters of processors 10 and 20 so as to increase the overall performance of portable communication device 50.

While certain features of the invention have been illustrated and described herein, many modifications, substitutions, changes, and equivalents will now occur to those skilled in the art. It is, therefore, to be understood that the appended
5 claims are intended to cover all such modifications and changes as fall within the true spirit of the invention.